arXiv:1609.09646v1 [math.NA] 30 Sep 2016

Nonlinear solution techniques for solving a Monge-Ampère equation for redistribution of a mesh

P.A. Browne[1], J. Prettyman[2,3], H. Weller[1], T. Pryer[2], and J. Van lent[4]

[1]Department of Meteorology, University of Reading, UK
[3]National Physical Laboratory, UK
[2]Department of Mathematics, University of Reading, UK
[4]FET – Engineering, Design and Mathematics, University of the West of England, UK
[*]Correspondence to p.browne@reading.ac.uk

October 3, 2016

Abstract

A Monge-Ampère equation arises when seeking an optimally transported mesh that equidistributes a given monitor function in Cartesian space. This Monge-Ampère equation is a fully nonlinear PDE, with a source term that is a function of the gradient of the solution. This nonlinear source term is an additional computational challenge that has received little attention from Monge-Ampère applications in other fields. There are two major components needed to find a solution to the Monge-Ampère equation: a spatial discretisation and an algorithm to find a solution of the resulting nonlinear algebraic equations. There have been a number of different approaches proposed in the literature to solve the Monge-Ampère equation but none of which perform consistent comparisons across both algorithmic and discretisation differences. In this study we explore different algorithmic methods for the Monge-Ampère equation all within the context of a finite volume spatial discretisation. We introduce a new linearisation of the Monge-Ampère equation that neglects the nonlinearities arising from the source term and show that it leads to a method that is fast, robust and free of tuning parameters. We present numerical experiments that show methods based on this linearisation of the Monge-Ampère equation are more computationally efficient than those that rely on other techniques such as a parabolic relaxation. Further, the equations resulting from a full linearisation of the Monge-Ampère equation, equivalent to using Newton's method, can be seen as analogous to an advection-diffusion equation. This allows many tools that exist for computational fluid dynamics to be re-factored easily to solve the Monge-Ampère equation. The robustness and efficiency of the newly introduced method gives hope that an adaptive solver for geophysical flows using mesh redistribution can be computationally feasible in the near future.

Graphical abstract

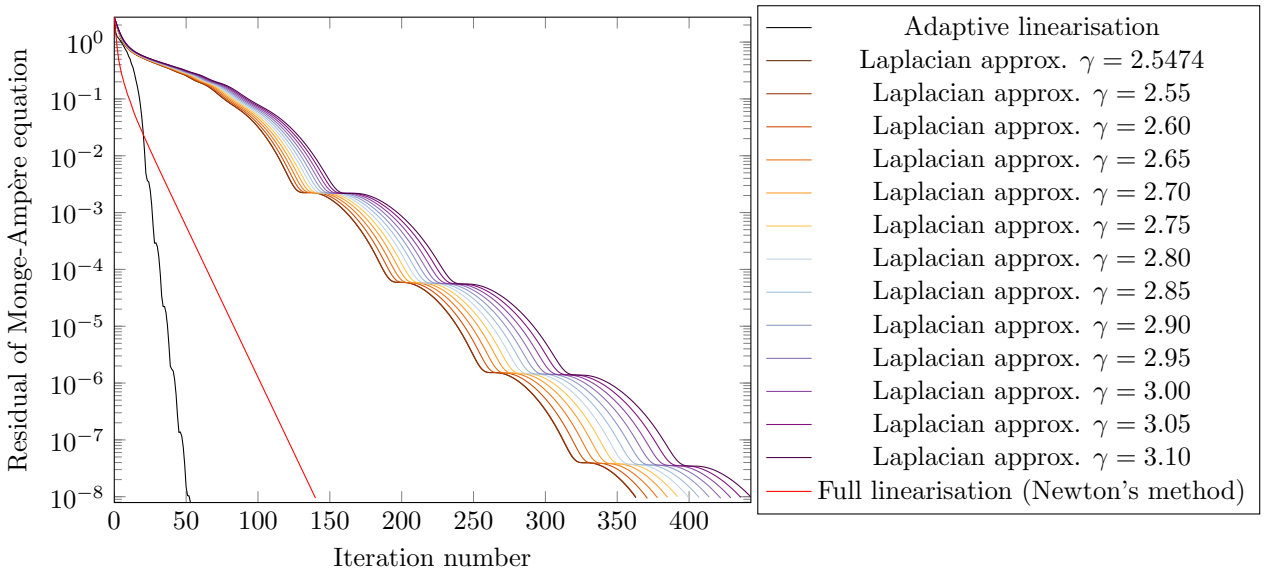

Keywords: Monge-Ampère equation, mesh redisribution, r-adaptivity, fixed point iterations, optimal transport, Newton's method

1 Introduction

Variable resolution meshes can be advantageous for the numerical solution of PDEs when variations or sensitivity to errors are greater in some areas than others. This occurs, for example, for the numerical prediction of tropical cyclones or for regional weather forecasts [Wang, 2001; Piani et al., 2000]. It may be advantageous that the mesh vary through time tracking atmospheric fronts [Kühnlein et al., 2012] or tsunamis [Harig et al., 2008]. Adding (or subtracting) mesh points where more (or fewer) points are required is known as *h-adaptivity*. With this type of adaptivity, the connectivity of the mesh (and possibly the total number of points defining the mesh) can change through time [Kimura et al., 2013].

R-adaptivity (mesh redistribution) involves keeping the mesh connectivity fixed but moving the mesh points. R-adaptivity retains fixed data structures associated with the mesh, need not entail mapping the solution from one mesh to another, may not lead to load balancing problems on parallel computers, can be easily incorporated into legacy code and can be designed to give meshes that vary smoothly in space and time. R-adaptivity can lead to smoothly graded meshes [Budd et al., 2009; Cao et al., 2003], which are may help to alleviate wave reflections or other errors associated with abrupt changes in resolution [Guba et al., 2014; Long and Thuburn, 2011; Vichnevetsky, 1987].

Optimal transport is a good technique for mesh-redistribution because it guarantees to find a mesh which is equidistributed with respect to a monitor function that is not tangled [Budd and Williams, 2009; Weller et al., 2016]. We aim to find a technique for finding optimally transported meshes which is robust, free of tunable parameters and fast enough to be used in numerical weather prediction when the mesh is being moved every time-step. In this article we will compare some of the existing solution techniques and introduce a new, faster and robust method.

A number of spatial discretistations and algorithms for solving non-linear equations have been introduced and used for solving the Monge-Ampère equation [e.g. Budd and Williams, 2009; Chacón et al., 2011; Weller et al., 2016] but some of these have free parameters and it is not clear which of these might be fast and robust enough for frequent adaptations during numerical weather prediction. In this article we shall use the finite volume spatial discretisation and perform a direct comparison of a number of existing and new algorithms for solving the Monge-Ampère equation in the context of mesh redistribution. Section 2 introduces the mathematical background of using the Monge-Ampère equation to find an optimally transported mesh as well as summarising the literature on solving the Monge-Ampère equation. In Section 3 we describe existing algorithms for solving the Monge-Ampère equation, and introduce new algorithms based on linearisations of different terms of the equation. In Section 4 we describe numerical experiments and diagnostics that we use to test the various algorithms. Finally, in Section 5 we draw conclusions about the different algorithmic methods for solving the Monge-Ampère equation for mesh redistribution.

2 Mathematical background

This section describes the theory of optimal transport for mesh redistribution and surveys the various approaches that have been taken in the literature to solve the resulting nonlinear problem.

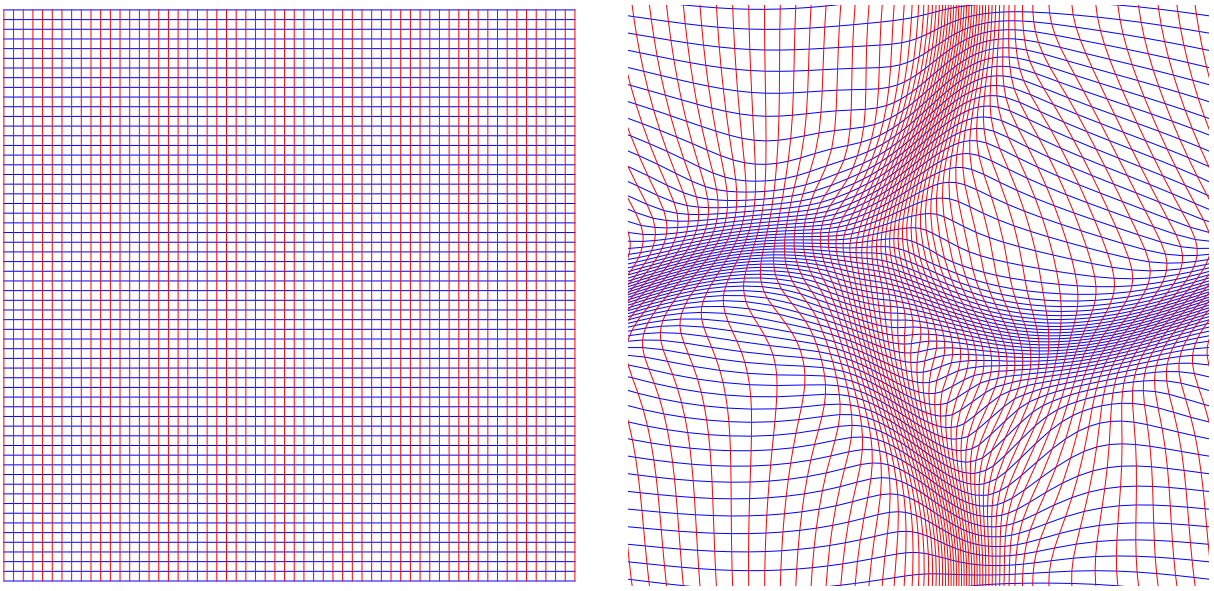

(a) Original computational mesh $\mathcal{T}_c$

(b) Adapted physical mesh $\mathcal{T}_p$

Figure 1: The action of a map $\boldsymbol{f}$ takes $\mathcal{T}_c$ to $\mathcal{T}_p$

Moving meshes, or r-adaptivity, is concerned with relocating the points of a mesh while keeping the mesh connectivity fixed. That is, an original mesh is given, $\mathcal{T}_c$, and it is transformed under a map $\boldsymbol{f}$ to $\mathcal{T}_p$ such that $\mathcal{T}_p = \boldsymbol{f}(\mathcal{T}_c)$. Figure 1 shows an example of these two meshes. The map, $\boldsymbol{f}$, needs to move points to regions which require higher resolution whilst retaining certain desirable properties of the original mesh, such as skewness, orthogonality or convexity.

2.1 Monitor function equidistribution

To help with notation, we introduce the computational and physical spaces, Ω_c and Ω_p, respectively. We have $\mathcal{T}_c \in \Omega_c$ and $\mathcal{T}_p \in \Omega_p$ and we write $\boldsymbol{\xi} \in \mathcal{T}_c$ and $\boldsymbol{x} \in \mathcal{T}_p$ to distinguish between elements of the two spaces.

To control the mesh adaptivity we assume that we have a scalar valued *monitor function* $m(\boldsymbol{x}) > 0$ given *a priori* which is large in parts of the domain which require a dense mesh and small in the parts of the domain which can be rarefied. We shall try to *equidistribute* this monitor function under the desired map.

That is, for each arbitrary volume in the physical mesh, the total monitor function contained within that volume is proportional to the volume of the inverse image of the volume, i.e. the corresponding volume in the original computational mesh. This can be rigorously written [Budd et al., 2009; Huang and Russell, 2011; Weller et al., 2016] as the equidistribution equation

$$m(\boldsymbol{x})|J(\boldsymbol{\xi})| = c, \tag{1}$$

where J is the Jacobian of the map $\boldsymbol{f}$ and c is a constant. If Ω_c and Ω_p are one-dimensional intervals then the solution of (1) uniquely defines the map $\boldsymbol{f}$. In higher dimensional problems (such as we consider in this article) the equidistribution equation is not enough to uniquely define $\boldsymbol{f}$; for example in a problem involving symmetries, $\boldsymbol{f}$ could be rotated and still satisfy equidistribution.

2.2 Optimally transported meshes

In order to arrive at a well-posed problem to find the adapted mesh, we must choose a method of regularising the equidistribution equation. There are many possible ways of doing so, such as specifying local conditions for the mesh. However, we note that the original computational mesh is likely to have desirable qualities, such as orthogonality, regularity etc. Hence it appears reasonable to seek an equidistributed mesh that is as close as possible to the computational mesh. In this sense we want to find an optimal map $\boldsymbol{f}^*$ such that

$$\boldsymbol{f}^* = \underset{\boldsymbol{f}:\Omega_c\to\Omega_p}{\arg\min} \int_{\Omega_c} |\boldsymbol{\xi} - \boldsymbol{f}(\boldsymbol{\xi})|^2 \, \mathrm{d}\boldsymbol{\xi}. \tag{2}$$

Once this map has been found, we call $\boldsymbol{f}^*$ an *optimal transport map* and the resulting mesh $\mathcal{T}_p = \boldsymbol{f}^*(\mathcal{T}_c)$ an *optimally transported mesh.*

2.3 The Monge-Ampère equation

A remarkable result from optimal transport theory is that the optimal mapping $\boldsymbol{f}^*$ is unique and can be written as the gradient of a convex potential [Brenier, 1991]. Hence

$$\boldsymbol{x} = \boldsymbol{f}^*(\boldsymbol{\xi}) = \nabla P(\boldsymbol{\xi}) \tag{3}$$

where P is a convex mesh potential. As P is convex it is easy to show that an optimally transported mesh does not exhibit tangling [Budd and Williams, 2009]. Substituting (3) into (1) we arrive at a Monge-Ampère equation:

$$m(\nabla P)|H(P)| = c \tag{4}$$

where $H(P)$ is the Hessian of the potential P, or equivalently the Jacobian of the map $\boldsymbol{f}^*$. Writing

$$\phi(\boldsymbol{\xi}) = P - \frac{|\boldsymbol{\xi}|^2}{2} \tag{5}$$

we arrive at the equation

$$m(\boldsymbol{\xi} + \nabla\phi)|I + H(\phi)| = c \tag{6}$$

which is a fully nonlinear elliptic PDE in the new mesh potential ϕ. The desired physical mesh is therefore given as

$$\boldsymbol{x} = \boldsymbol{\xi} + \nabla\phi. \tag{7}$$

As $m > 0$, we can re-write the Monge-Ampère equation (6) as

$$|I + H(\phi)| = \frac{c}{m(\boldsymbol{x})}. \tag{8}$$

To numerically solve the Monge-Ampère equation requires two components: a spatial discretisation to get a discrete representation of the equation and an algorithm which finds the solution to the resulting system of nonlinear algebraic equations.

2.3.1 Spatial Discretisations

Spatial discretisation of the Monge-Ampère equation has mostly focused on the discretisation of $|H(\phi)|$ [e.g. Froese and Oberman, 2011]. Oberman [2008] developed wide stencil finite difference techniques based on the Barles and Souganidis [1991] framework to ensure discrete monotonicity guaranteeing convergence of the numerical solution to the unique convex viscosity solution. Mixed finite-element methods have been used by Dean and Glowinski [2006b]; Feng and Neilan [2009]; Lakkis and Pryer [2013]; Neilan [2014]. Finite differences have been used extensively in solving the Monge-Ampère equation for mesh redistribution problems [e.g. Delzanno et al., 2008; Sulman et al., 2011]. Budd et al. [2009]; Browne et al. [2014] used simple finite differences with filtering of the right hand side and smoothing of the Hessian in order to ensure that a convex solution is found. Weller et al. [2016] used finite volumes to solve the Monge-Ampère equation on a plane and on the sphere and explored a number of techniques for solving the $m(\boldsymbol{\xi} + \nabla\phi)$ term on the right hand side of eqn (8) while Froese [2012] used wide stencil finite differences to discretise $\nabla\phi$ on the right hand side of the Monge-Ampère equation and Saumier et al. [2015] experimented with finite differences and a spectral method. Feng and Jensen [2016] give a reformulation of the Monge-Ampère problem as a Hamilton-Jacobi-Bellman equation which removes the constraint of convexity on the solution as a by-product of this reformulation.

In this article we shall consider using only the finite volume technique similarly to Weller et al. [2016] in order to compare various differences in algorithmic processes for solving the Monge-Ampère equation.

2.3.2 Non-linear Equation Solution Algorithms

The Monge-Ampère equation can be posed as the nonlinear PDE

$$|H(\phi)| = g.$$

The majority of the literature on solving this Monge-Ampère equation is devoted to the case where the source term g is a given function of spatial location on the computational mesh [e.g. Dean and Glowinski, 2006b; Feng and Neilan, 2009; Oberman, 2008]. For the mesh redistribution, this source term becomes a function of the solution ϕ itself (i.e. $g = g(\nabla\phi)$) since m is defined in physical space. Techniques for treating this nonlinearity in the source term have been investigated by Froese [2012] and Weller et al. [2016]. This nonlinearity also occurs in the case of prescribed Gauss curvature [see for example Pryer, 2012].

Benamou et al. [2010] give a concise overview of the algorithmic methods for solving the Monge-Ampère equation when g is not directly a function of the solution. Dean and Glowinski [2006a,b] use an Augmented Lagrangian method to the Monge-Ampère equation after it has been reformulated as a saddle point problem. Dean and Glowinski [2006b] use a nonlinear least-squares solution to minimise an appropriate functional of the solution. Benamou et al. [2010] use a fixed point method derived by taking a Poisson approximation to the Monge-Ampère equation. A related fixed point method using a Poisson approximation was introduced by Weller et al. [2016] for solving an equation of Monge-Ampère type for mesh redistribution on a spherical manifold.

For mesh redistribution via optimal transport in Cartesian geometry, Delzanno et al. [2008] and Chacón et al. [2011] use a multigrid-preconditioned inexact Newton-Krylov method with damping. A parabolic relaxation to the Monge-Ampère equation was proposed by Budd and Williams [2009] and has been used successfully in a number of studies ([Budd et al., 2013; Browne et al., 2014]), with other forms of parabolic relaxation investigated by Sulman et al. [2011].

3 Details of techniques for solving the Monge-Ampère equation

In this section we shall describe in detail algorithms for solving the Monge-Ampère equation in the context of mesh redistribution where the source term is a function of the gradient of the solution. In particular, we will introduce a linearisation about the current approximate solution which will lead to a fixed point method with no free parameters. We will then describe the spatial discretisation.The essential computational components of the algorithms are summarised in Table 1.

3.1 Fixed point iterations

In order to find a solution to the Monge-Ampère equation, Weller et al. [2016] introduced a fixed point method, which we describe here.

Firstly, as $m > 0$, we can re-write the Monge-Ampère equation (6) as

$$|I + H(\phi)| - \frac{c}{m(\boldsymbol{x})} = 0 \tag{9}$$

Let ϕ^n denote the approximation to the solution at iteration n. Taking a Taylor's series expansion of $|I + H(\phi^{n+1})|$ about $\phi^0 = 0$ we see

$$|I + H(\phi^{n+1})| = 1 + \nabla^2\phi^{n+1} + \mathcal{N}(\phi^{n+1}) \tag{10}$$

where $\mathcal{N}(\phi^{n+1})$ are higher order, nonlinear terms. In 2D we have that the nonlinear terms are precisely $\mathcal{N}(\phi^{n+1}) = |H(\phi^{n+1})|$.

We want to solve

$$|I + H(\phi^{n+1})| = \frac{c}{m(\boldsymbol{x}^{n+1})}. \tag{11}$$

Substituting (10) into (11), and subtracting $|I + H(\phi^n)|$ we obtain

$$1 + \nabla^2\phi^{n+1} + \mathcal{N}(\phi^{n+1}) - \mathcal{N}(\phi^n) = 1 + \nabla^2\phi^n - |I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^{n+1})}. \tag{12}$$

Dropping the change in nonlinear terms between iterations n and $n + 1$ gives a fixed point iteration in ϕ^n such that, given $\boldsymbol{x}^0$ and ϕ^0,

$$1 + \nabla^2\phi^{n+1} = 1 + \nabla^2\phi^n - |I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}, \qquad \forall n \in \mathbb{N}. \tag{13}$$

It is easy to see that a fixed point of this iteration will solve the fully nonlinear Monge-Ampère equation. For this fixed point iteration to be stable we require that the nonlinear terms we have omitted, $\mathcal{N}(\phi^{n+1}) - \mathcal{N}(\phi^n)$, are small. Equivalently, $|H(\phi^{n+1})| \approx |H(\phi^n)|$. Note that these are nonlinear functions of variables the same order as ϕ.

As Weller et al. [2016] showed (and can be seen in Section 4) this fixed point iteration is not always convergent. Under-relaxation can force this equation to converge. That is, introducing a scalar constant γ, we can weight the linear terms in (13) such that

$$\gamma\nabla^2\phi^{n+1} = \gamma\nabla^2\phi^n - |I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}, \qquad \forall n \in \mathbb{N}. \tag{14}$$

For large γ there is less weight given to the error in the nonlinear Monge-Ampère equation and hence only a small update is made to the approximate solution ϕ. Vice versa, the smaller γ, the more the nonlinear problem is treated directly, and hence the large the update from each fixed point iteration. We have no theoretical basis on which to find the optimal value of γ which, as will be seen in Section 4, is dependent on the monitor function in question. In order to have a method with no free parameters, we introduce a different linearisation of the Monge-Ampère equation.

3.2 A fixed point method with an adaptive linearisation

Instead of linearising $|I+H(\phi^{n+1})|$ about ϕ^0, we can instead linearise about a current iterate ϕ^n. Writing $\phi^{n+1} = \phi^n + \varepsilon\psi$ it can be shown that

$$|I + H(\phi^{n+1})| = |I + H(\phi^n)| + \nabla \cdot A^n\nabla\varepsilon\psi + \mathcal{N}(\varepsilon\psi), \tag{15}$$

where A^n is the matrix of cofactors of $I + H(\phi)$ and $\mathcal{N}$ is some nonlinear function. In 2D

$$A^n = \begin{bmatrix} 1 + \phi^n_{yy} & -\phi^n_{xy} \\ -\phi^n_{xy} & 1 + \phi^n_{xx} \end{bmatrix} \tag{16}$$

and $\mathcal{N}(\varepsilon\psi) = \varepsilon^2|H(\phi)|$. In 3D, a more involved computation can show $\mathcal{N}(\varepsilon\psi) = \varepsilon^3\tilde{\mathcal{N}}(\psi)$ and

$$A^n = \begin{bmatrix} 1 + \phi^n_{yy} + \phi^n_{zz} + \phi^n_{yy}\phi^n_{zz} - \phi^n_{yz}\phi^n_{yz} & -\phi^n_{xy} - \phi^n_{xy}\phi^n_{zz} + \phi^n_{xz}\phi^n_{yz} & -\phi^n_{xz} - \phi^n_{xz}\phi^n_{yy} + \phi^n_{xy}\phi^n_{yz} \\ -\phi^n_{xy} - \phi^n_{xy}\phi^n_{zz} + \phi^n_{xz}\phi^n_{yz} & 1 + \phi^n_{xx} + \phi^n_{zz} + \phi^n_{xx}\phi^n_{zz} - \phi^n_{xz}\phi^n_{xz} & -\phi^n_{yz} - \phi^n_{xx}\phi^n_{yz} + \phi^n_{xy}\phi^n_{xz} \\ -\phi^n_{xz} - \phi^n_{xz}\phi^n_{yy} + \phi^n_{xy}\phi^n_{yz} & -\phi^n_{yz} - \phi^n_{xx}\phi^n_{yz} + \phi^n_{xy}\phi^n_{xz} & 1 + \phi^n_{xx} + \phi^n_{yy} + \phi^n_{xx}\phi^n_{yy} - \phi^n_{xy}\phi^n_{xy} \end{bmatrix}. \tag{17}$$

Hence defining A^n as above and substituting (15) into the Monge-Ampère equation (11) we obtain

$$|I + H(\phi^n)| + \nabla \cdot (A^n\nabla\varepsilon\psi) + \mathcal{N}(\varepsilon\psi) = \frac{c}{m(\boldsymbol{x}^{n+1})}. \tag{18}$$

Dropping the terms proportional to ε^d where d is the dimension of the space gives a fixed point iteration in ϕ^n such that, given $\boldsymbol{x}^0$ and ϕ^0,

$$\nabla \cdot (A^n\nabla\varepsilon\psi) = -|I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}, \qquad \forall n \in \mathbb{N}. \tag{19}$$

As in (13), a fixed point of (19) solves the Monge-Ampère equation. The nonlinear term we have omitted in (13) is $\mathcal{N}(\varepsilon\psi) = \varepsilon^d|H(\psi)|$. When ε is small (i.e. $\phi^{n+1} \approx \phi^n$) the nonlinear terms will be smaller than

those omitted in the fixed point method (13). Note that at $\phi^0 = 0$, the (19) reduces to the initial fixed point method as $H(0) = 0 \implies A^0 = 0$.

The iterative method for solving the Monge-Ampère equation given in (19) can have numerical difficulties when the discretised, matrix equation becomes indefinite. In this case, the ellipticity property of the original Monge-Ampère equation is lost and the nonconvex solutions to the Monge-Ampère equation can be generated which lead to tangled meshes and numerical divergence. The indefiniteness of (19) is caused by the matrix A^n being numerically indefinite. Therefore, we can modify (19) to maintain ellipticity by

$$\nabla \cdot (B^n \nabla \varepsilon \psi) = -|I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}, \qquad \forall n \in \mathbb{N}, \tag{20}$$

where

$$B^n = A^n + \gamma I \tag{21}$$

and γ is defined as

$$\gamma := \begin{cases} 0 & \text{if} \quad \min \sigma[A^n] > 0 \\ \epsilon - \min \sigma[A^n] & \text{if} \quad \min \sigma[A^n] \leq 0. \end{cases} \tag{22}$$

The constant $\epsilon > 0$ is chosen to avoid round-off errors (we have taken $\epsilon = 10^{-5}$), and $\sigma[A^n]$ refers to the spectrum of A^n. This process simply shifts the eigenvalues of the matrix A^n so that they remain positive. One can check that this choice of γ ensures that $B^n \succ 0$. We shall refer to the algorithm given in (20) as the Adaptive Fixed Point method (AFP).

There are two key remarks about about (20): firstly there are no free parameters in the method and, secondly, the scalar γ can be spatially varying. The lack of free parameters makes this very attractive as a robust method for adapting a mesh. Such robust methods will be required for the future coupling of optimally transported meshes into operational forecasting models where the monitor function that can vary in time as well as space is not known *a priori*. Instead of using one global γ (as for the FP method in Section 3.1) to regularise the global system, this choice of regularisation shifts only the eigenvalues of the local matrices A^n in the parts of the domain where they are numerically indefinite.

A further remark about (20) is that is can be discretised in space using a finite volume, finite element or finite difference method and solved implicitly for ϕ^{n+1}. It can be seen as a nonconstant tensor coefficient Poisson equation. This type of equation is particularly difficult to solve with a spectral method as it is nonseparable [Boyd, 2013].

3.3 Newton's method for the Monge-Ampère equation

Starting from (11), we wish to solve

$$|I + H(\phi^{n+1})| = \frac{c}{m(\boldsymbol{x}^{n+1})}. \tag{23}$$

Numerical solution using Newton's method involves linearising not only $|I+H(\phi)|$, but also the nonlinear right hand side $\frac{c}{m}(\phi)$.

Linearising the right hand side $\frac{c}{m(\boldsymbol{x}^{n+1})}$ about $\boldsymbol{x}^n$ gives

$$\frac{c}{m}(\boldsymbol{x}^{n+1}) = \frac{c}{m}(\boldsymbol{x}^n) + \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot(\boldsymbol{x}^{n+1}-\boldsymbol{x}^n). \tag{24}$$

Recall $\boldsymbol{x}^{n+1} = \boldsymbol{\xi} + \nabla\phi^{n+1}$ and $\boldsymbol{x}^n = \boldsymbol{\xi} + \nabla\phi^n$, thus $\boldsymbol{x}^{n+1} - \boldsymbol{x}^n = \nabla\phi^{n+1} - \nabla\phi^n$. Also we write $\phi^{n+1} = \phi^n + \varepsilon\psi$. Hence

$$\frac{c}{m}(\boldsymbol{x}^{n+1}) = \frac{c}{m}(\boldsymbol{x}^n) + \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot(\nabla\phi^{n+1}-\nabla\phi^n) \tag{25}$$

$$= \frac{c}{m}(\boldsymbol{x}^n) + \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot\nabla(\varepsilon\psi). \tag{26}$$

This can be incorporated into (18) to obtain the linearisation of the full Monge-Ampère equation. Note that (26) contains gradients in two different spaces: ∇, the gradient on the computational mesh, and $\nabla_{\boldsymbol{x}}$, the gradient on the physical mesh.

Substituting (26) and (15) into (23) gives:

$$|I + H(\phi^n)| + \nabla\cdot(B^n\nabla\varepsilon\psi) = \frac{c}{m(\boldsymbol{x}^n)} + \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot\nabla(\varepsilon\psi). \tag{27}$$

Rearranging the terms leads to

$$\nabla\cdot(B^n\nabla\varepsilon\psi) - \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot\nabla(\varepsilon\psi) + |I + H(\phi^n)| - \frac{c}{m(\boldsymbol{x}^n)} = 0. \tag{28}$$

Iterating equation (28) is Newton's method for solving the Monge-Ampère equation. Written in this form it can be seen to have a physical interpretation: it is an advection–diffusion equation for $\varepsilon\psi$ with tensorial diffusion coefficient B^n, advection velocity $\nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)$ and source term $-|I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}$.

As we can write Newton's method for solving the Monge-Ampère equation as an advection–diffusion equation, we can use the tools and knowledge from computational fluid dynamics to find a numerical solution. This will be discussed further in Section 3.5.1.

3.4 The Parabolic Monge-Ampère method

The parabolic Monge-Ampère (PMA) method for solving the Monge-Ampère equation was introduced by Budd and Williams [2009]. They consider a pseudo-time equation

$$(I - \gamma\nabla^2)\phi_\tau = [m(\boldsymbol{x})|I + H(\phi)|]^{\frac{1}{d}} \tag{29}$$

where γ is a smoothing coefficient (constant over space), ϕ_τ is the pseudo-time derivative of the potential ϕ and d is the dimension of the space in which the Monge-Ampère equation is being solved. Budd and

Table 1: Comparison of algorithmic methods for solving the Monge-Ampère equation

Method	Inversion	Free parameters	Advection term
Fixed point (FP)	∇^2	γ	None
Adaptive fixed point (AFP)	$\nabla \cdot B^n \nabla$	None	None
Newton's method	$\nabla \cdot B^n \nabla$	δ	Yes
Parabolic relaxation (PMA)	$I - \gamma \nabla^2$	γ, δt	None

Williams [2009] showed that as $\tau \to \infty$, $\nabla\phi \to \nabla P$ where ∇P solves the original Monge-Ampère equation. Other parabolic relaxations of the Monge-Ampère equation have been investigated in the literature. For example, Sulman et al. [2011] use a logarithmic form of relaxation instead of the power law in (29).

In the PMA method, only the gradient of the potential is of interest. From a numerical perspective, it is useful to note that the potential is always increasing: $m > 0$ and $|H| > 0$ as H is symmetric positive definite. Hence this potential can be modified by the addition of any constant without affecting the PMA method's convergence to the solution of the original Monge-Ampère equation.

We can choose this constant to be

$$c = -\int_\Omega (I - \gamma\nabla^2)^{-1} \left[m(\boldsymbol{x})|I + H(\phi)|\right]^{\frac{1}{d}} \tag{30}$$

so that

$$\int_\Omega \phi_\tau = 0 \implies \int_\Omega \phi^n = \int_\Omega \phi^0 \quad \forall n \in \mathbb{N}. \tag{31}$$

With this choice of constant, PMA can be written as a fixed point method in ϕ^n

$$(I - \gamma\nabla^2)\phi^{n+1} = (I - \gamma\nabla^2)\phi^n + \delta t \left[m(\boldsymbol{x}^n)|I + H(\phi^n)|\right]^{\frac{1}{d}} + c. \tag{32}$$

Note that PMA has 2 separate parameters to choose in order to solve the Monge-Ampère equation: γ and δt. We consider the PMA method here as a benchmark with which to compare the fixed point and adaptive linearisation methods due to the considerable literature using PMA for mesh adaptation (e.g. [Budd and Williams, 2009; Budd et al., 2013; Browne et al., 2014] etc).

The algorithmic methods which we consider in this article are summarised in Table 1 that shows what matrix equation is solved at each timestep and what free parameters each method needs to set.

3.5 Spatial Discretisation

We use a finite volume technique to discretise the Monge-Ampère equation in space using OpenFOAM [www.openfoam.org, 2015], following Weller et al. [2016]. The discretisation assumes that all finite volume

cells are three dimensional with two dimensional faces between cells. The test cases in this paper use one layer of cells to represent the two dimensional domain. The prognostic variable ϕ is stored at the centre of each cell. To calculate the value of the Laplacian in cell i, $(\nabla^2\phi)_i$ we use the Divergence Theorem, and write it as a sum over each face f of the cell. Hence

$$(\nabla^2\phi)_i \approx \frac{1}{V_i}\sum_{f\in i}|\boldsymbol{S}_f|\frac{\phi_{i_f}-\phi_i}{|\boldsymbol{d}_f|} \tag{33}$$

where V_i is the volume of cell i, $\boldsymbol{S}_f$ is the vector normal to the face f with magnitude $|\boldsymbol{S}_f|$ equal to the area of the face f, i_f refers to the cell connected to cell i via the face f and $|\boldsymbol{d}_f|$ is the distance between the centre of cell i_f and the centre of cell i. The notation $f \in i$ refers to a face f of cell i.

The gradient of ϕ at cell centres. $\nabla_c\phi$ is computed using the divergence theorem such that

$$(\nabla_c\phi)_i = \frac{1}{V_i}\sum_{f\in i}\phi_f\boldsymbol{S}_f \tag{34}$$

where ϕ_f is the value of ϕ linearly interpolated onto face f. The gradient of ϕ can be computed in the direction normal to a face f by

$$\nabla_{nf}\phi = \frac{\phi_{i_f}-\phi_i}{|\boldsymbol{d}_f|}. \tag{35}$$

To calculate the value of the Hessian in cell i, $H_i(\phi)$ we first compute $\nabla_c\phi$ which is valid on the cell centres and linearly interpolate this onto the faces to give $\widetilde{\nabla_f\phi}$. The normal component of the full gradient on the face, $\nabla_f\phi$, is corrected such that

$$\nabla_f\phi = \widetilde{\nabla_f\phi} + \left(\nabla_{nf}\phi - (\widetilde{\nabla_f\phi}\cdot\hat{\boldsymbol{S}}_f)\right)\hat{\boldsymbol{S}}_f \tag{36}$$

where $\hat{\boldsymbol{S}}_f$ is the unit vector normal to face f.

Then we compute the gradient of this which is valid at the cell centres using the divergence theorem.

$$H_i(\phi) = \nabla\nabla\phi = \frac{1}{V_i}\sum_{f\in i}(\nabla_f\phi)\boldsymbol{S}_f \tag{37}$$

We store the mesh points at the cell corners. Therefore to update the mesh we need to have a gradient $\nabla\phi$ valid on cell corners. To do this, we reconstruct the gradient at the corners from $\nabla_{nf}\phi$ which, in the case of a uniform grid, is the average the gradient on each face that contains the cell corner in question.

3.5.1 Spatial discretisation of the advection term for Newton's method

Recall that in (28) we wrote Newton's method for solving the Monge-Ampère equation as an advection–diffusion equation. This will be solved using existing functions in OpenFOAM which can solve advection-diffusion equations implicitly if the advection term is written in conservative form. To go from advective

to conservative form, we note the following vector calculus identity:

$$\nabla_{\boldsymbol{x}} a \cdot \nabla_{\boldsymbol{y}} b = \nabla_{\boldsymbol{y}} \cdot ((\nabla_{\boldsymbol{x}} a) b) - b \nabla_{\boldsymbol{y}} \cdot \nabla_{\boldsymbol{x}} a \tag{38}$$

Using the identity (38), the advection term in (28) becomes

$$\nabla_{\boldsymbol{x}} \left(\frac{c}{m}(\boldsymbol{x}^n) \right) \cdot \nabla \psi = \nabla \cdot (\nabla_{\boldsymbol{x}} \left(\frac{c}{m}(\boldsymbol{x}^n) \right) \varepsilon \psi) - \varepsilon \psi \nabla \cdot \nabla_{\boldsymbol{x}} \left(\frac{c}{m}(\boldsymbol{x}^n) \right)$$

and so (changing sign to follow the usual convention) Newton's method for the Monge-Ampère equation can be written

$$- \nabla \cdot (B^n \nabla \varepsilon \psi) + \nabla \cdot \left(\nabla_{\boldsymbol{x}} \left(\frac{c}{m}(\boldsymbol{x}^n) \right) \varepsilon \psi \right) - \varepsilon \psi \nabla_{\boldsymbol{x}} \cdot \nabla \left(\frac{c}{m}(\boldsymbol{x}^n) \right) - |I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)} = 0. \tag{39}$$

Equation (39) is what we actually compute with using the finite volume software OpenFOAM.

Note that $\nabla_x(\frac{c}{m})$ has to be computed on the physical mesh $\mathcal{T}_p$: when this mesh becomes non-orthogonal this can lead to errors if we were to compute the gradient using the normal directions to each cell as in (35). Instead, for this gradient calculation, we follow Weller [2014] and compute the gradient in the direction $\boldsymbol{d}_i$ that goes from cell centre to cell centre. Consider cell i of the physical mesh with faces indexed by f and neighbour cells indexed by N_f.

For each face, f, $\boldsymbol{d}_f$ is the vector from the centre of cell i, $\boldsymbol{x}_i$, to the centre of cell N_f, $\boldsymbol{x}_{N_f}$, i.e. $\boldsymbol{d}_f = \boldsymbol{x}_{N_f} - \boldsymbol{x}_i$. The gradient, $\nabla_x(\frac{c}{m})$ is calculated for cell i as:

$$\nabla_x \left(\frac{c}{m} \right) = \left(\sum_{f \in i} \boldsymbol{d}_f \boldsymbol{d}_f^T \right)^{-1} \sum_{f \in i} \boldsymbol{d}_f \left(\left(\frac{c}{m} \right)_{N_f} - \left(\frac{c}{m} \right)_i \right). \tag{40}$$

3.6 Solving the tensorial diffusion coefficient Poisson equation

In order to find the update for both the AFP method and Newton's method we must solve an equation of the form

$$\nabla \cdot B \nabla \Psi = b \tag{41}$$

which gives rise to a matrix equation

$$M\boldsymbol{x} = \boldsymbol{b}. \tag{42}$$

To compute the entries of the matrix M, we first split the gradient of Ψ into the components normal and tangential to cell faces such that

$$\nabla \Psi = \hat{\mathbf{S}}_f \nabla_{nf} \Psi + \left((\nabla \Psi)_f - \left((\nabla \Psi)_f \cdot \hat{\mathbf{S}}_f \right) \hat{\mathbf{S}}_f \right) \tag{43}$$

where $\mathbf{S}_f$ is the face area vector; the vector normal to each face with magnitude equal to the face area, $\hat{\mathbf{S}}_f$ is the unit normal vector to each face, $\nabla_{nf}\Psi$ is the gradient of Ψ calculated in direction $\hat{\mathbf{S}}_f$ using (35)

and $\nabla\Psi$ is the gradient of Ψ calculated at cell centres using the divergence theorem (similarly to (34)) and the $()_f$ notation means that it is then interpolated from cell centres to faces.

We treat implicitly the normal components of the gradient and explicitly the tangential components of the gradient. Substituting (43) into (41), using the linearity of the divergence operator and re-arranging so that implicitly calculated parts are on the left hand side of the equation gives the linear equation:

$$\nabla \cdot \left(B\hat{\mathbf{S}}_f \nabla_n \Psi \right) = b - \nabla \cdot \left(B \left((\nabla\Psi)_f - \left((\nabla\Psi)_f \cdot \hat{\mathbf{S}}_f \right) \hat{\mathbf{S}}_f \right) \right). \tag{44}$$

and solve this iteratively, updating the right hand side based on the current estimate of the solution. We have found that 3 iterations of this is sufficient to make the outer iteration process of the AFP and Newton method stable and robust. With fewer iterations the number of fixed point (or Newton) iterations to solve the Monge-Ampère equation increased, and with more the number of fixed point iterations did not notably change.

4 Numerical experiments

4.1 Test cases

Meshes are generated to equidistribute a static monitor function in two dimensions on the domain $\left[-\frac{1}{2}, \frac{1}{2}\right]^2$. We follow Budd et al. [2015] and Weller et al. [2016] in using a radially symmetric monitor function of the form

$$m(\boldsymbol{x}) = 1 + \alpha_1 \mathrm{sech}^2 \left(\alpha_2 \left(|\boldsymbol{x}|^2 - \alpha_3^2 \right) \right). \tag{45}$$

Using this form we define two different monitor functions: the *ring* function, where $\alpha_1 = 10$, $\alpha_2 = 200$ and $\alpha_3 = 0.25$; and the *bell* function, where $\alpha_1 = 50$, $\alpha_2 = 100$ and $\alpha_3 = 0$. We solve the Monge-Ampère equation with periodic boundary conditions in both the x and y directions.

For each test problem, the computational mesh $\mathcal{T}_c$ is a uniformly spaced quadrilateral grid as depicted in Figure 1a. The number of cells in each example will vary and be denoted $N \times N$.

4.2 Numerical linear algebra details

It is well known that the Laplace operator on a periodic domain has a zero eigenvalue corresponding to the constant eigenfunction. Therefore we must remove the kernel of both ∇^2 in the FP method and $\nabla \cdot (B^n \nabla)$ in the AFP method. We do so by setting a fixed reference value for one of the points in the solution – the choice of which point and what value do not influence the results.

For the Newton method, we have an added advection term in the implicit equation to be solved at each timestep. At each iteration it is an advection-diffusion equation with the Monge-Ampère equation

as a source term. Thus each iteration itself is hyperbolic. The boundary conditions of the hyperbolic part of the equation need to be set along every characteristic. This is not feasible as the characteristics change with every iteration. To have a well-posed numerical problem at each iteration we impose that the solution integrates to zero globally. This is done by solving the equation

$$\delta\varepsilon\psi + \nabla\cdot(B^n\nabla\varepsilon\psi) = \nabla_{\boldsymbol{x}}\left(\frac{c}{m}(\boldsymbol{x}^n)\right)\cdot\nabla(\varepsilon\psi) - |I + H(\phi^n)| + \frac{c}{m(\boldsymbol{x}^n)}, \tag{46}$$

where δ is a scalar with dimensions area^{-1}. In the 2D examples considered in this article we set $\delta = \frac{t10^{-4}}{\min_i V_i}$, where V_i is the volume of cell i in the computational mesh and t =1m is the thickness of the domain in the direction normal to the solution domain. An appropriate value to use in 3D has yet to be determined.

The linear solver used is geometric-algebraic multigrid (GAMG) with a symmetric Gauss-Seidel smoother. The solver tolerance was set to 10^{-12} with a relative tolerance of 10^{-8}. Details of this solver can be found in the documentation of OpenFOAM [www.openfoam.org, 2015].

To compute the eigenvalues that appear in (22) we use LAPACK [Anderson et al., 1999] – we found that the eigenvalue calculation internal to OpenFOAM which attempts to find roots of the characteristic polynomial of the matrix is not robust to round-off errors.

4.3 Diagnostics

We wish to compare the efficiency of the various solution techniques in solving the Monge-Ampère equation. CPU time is of course the overall goal for any numerical method, however this is not necessarily a robust measure as it will depend on both hardware and software implementation. Figure 2 shows, for each method, the CPU time plotted against iteration count. We will consider the total number of outer iterations as a measure of the efficiency of the methods, and leave optimization of the codes for future investigations.

To measure accuracy, we define an equidistribution measure following Browne et al. [2014]: given that (6) says that $m(\boldsymbol{x})|I + H(\phi)| =$ constant at all points in the domain, we look at the coefficient of variation of this quantity over all the points in the mesh. That is,

$$\varepsilon = \frac{\sqrt{\mathrm{Var}\{m(\boldsymbol{x})|I + H(\phi)|\}}}{\overline{m(\boldsymbol{x})|I + H(\phi)|}}. \tag{47}$$

Clearly when $\varepsilon \to 0$, $m(\boldsymbol{x})|I + H(\phi)|$ is constant and thus the Monge-Ampère equation is satisfied.

For all the test cases presented in this article, we iterate each method until the equidistribution $\varepsilon < 10^{-8}$. This is almost certainly an unnecessarily tight tolerance for mesh generation, however it will illustrate the convergence of each method to high accuracy.

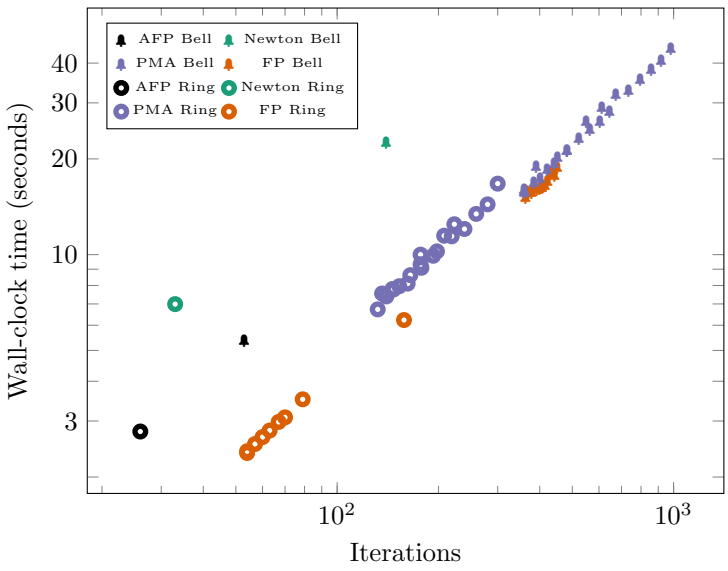

Figure 2: Iterations vs Wall-clock time for a number of different experiments on the Ring and Bell test cases. These are for a computational mesh of 60×60 cells, with the codes all running in serial. Both Newton's method and the AFP method are more costly per iteration due to the 3 solves at each iteration that we need in OpenFOAM to solve the linear system - this is a feature of our implementation and not necessarily the algorithms themselves. Newton's method is more costly per iteration than all the other methods due to the extra computations needed to calculate the advection term and is an inherent feature of that algorithm.

4.4 Results

Figure 3 shows the resulting meshes for both the ring and the bell test case, when using the AFP method with $N = 60$. At this resolution, the resulting meshes for all converged solution techniques appear identical to the eye. This is unsurprising given the unique solution to the optimal transport problem and the very tight tolerance to which we are solving the Monge-Ampère equation with each method, and the fact that they all use the same spatial discretisation.

It can be seen from Figures 4 and 5 that both the fixed point iterations and the PMA technique exhibit linear convergence. However the rate of this convergence is dependent on the choice of parameters used. The AFP method, with no free parameters also shows linear convergence but at a must faster rate. Newton's method is seen to have the best performance initially when the solution is far from converged. This is due to the advection term that accelerates the solution by moving the mesh points up the gradient of the monitor function.

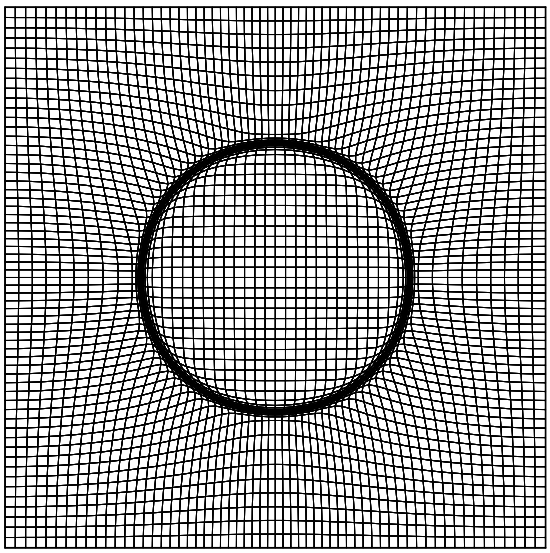

(a) Grid generated with AFP method for the ring test case

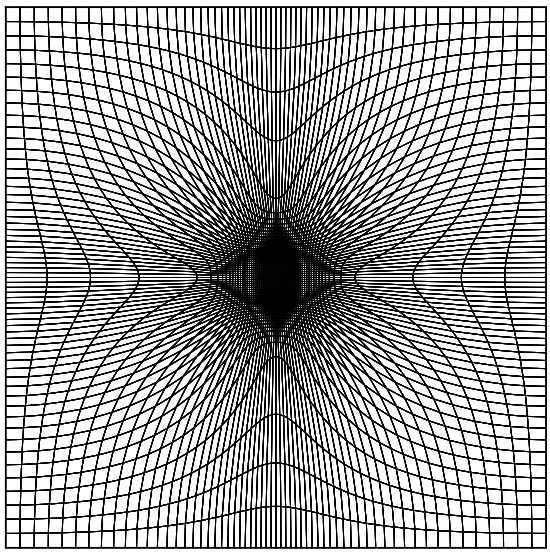

(b) Grid generated with AFP method for the bell test case

Figure 3: Resulting grids for the two different test cases

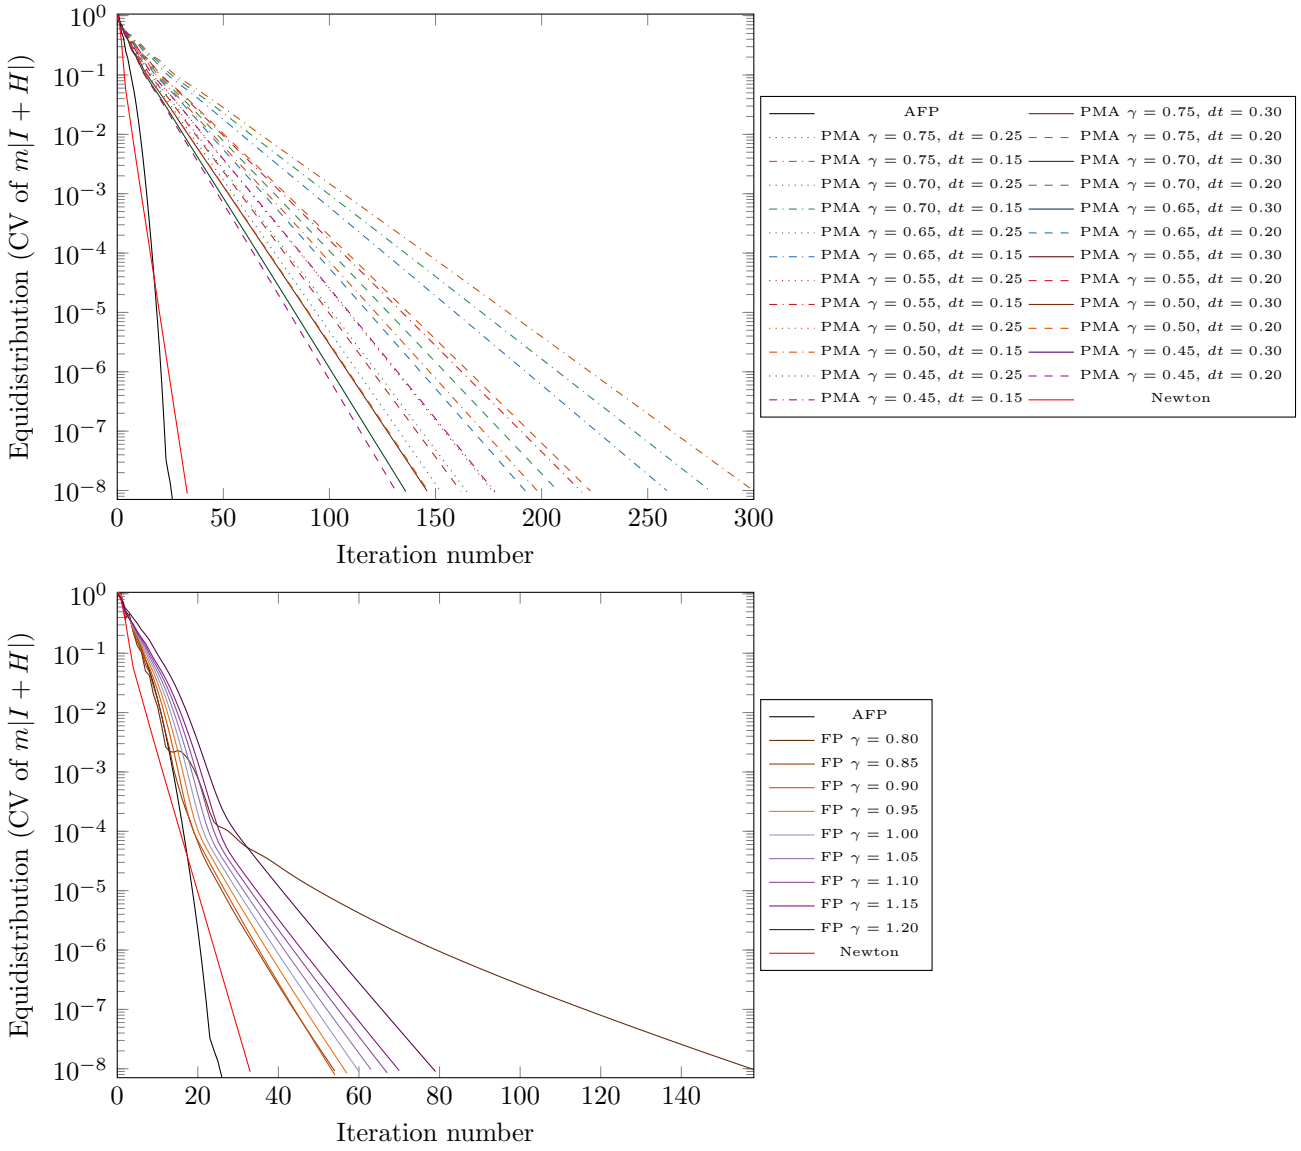

Figure 4: Plots of equidistribution against iteration number for the different methods when applied to the ring test case with an initial 60×60 mesh

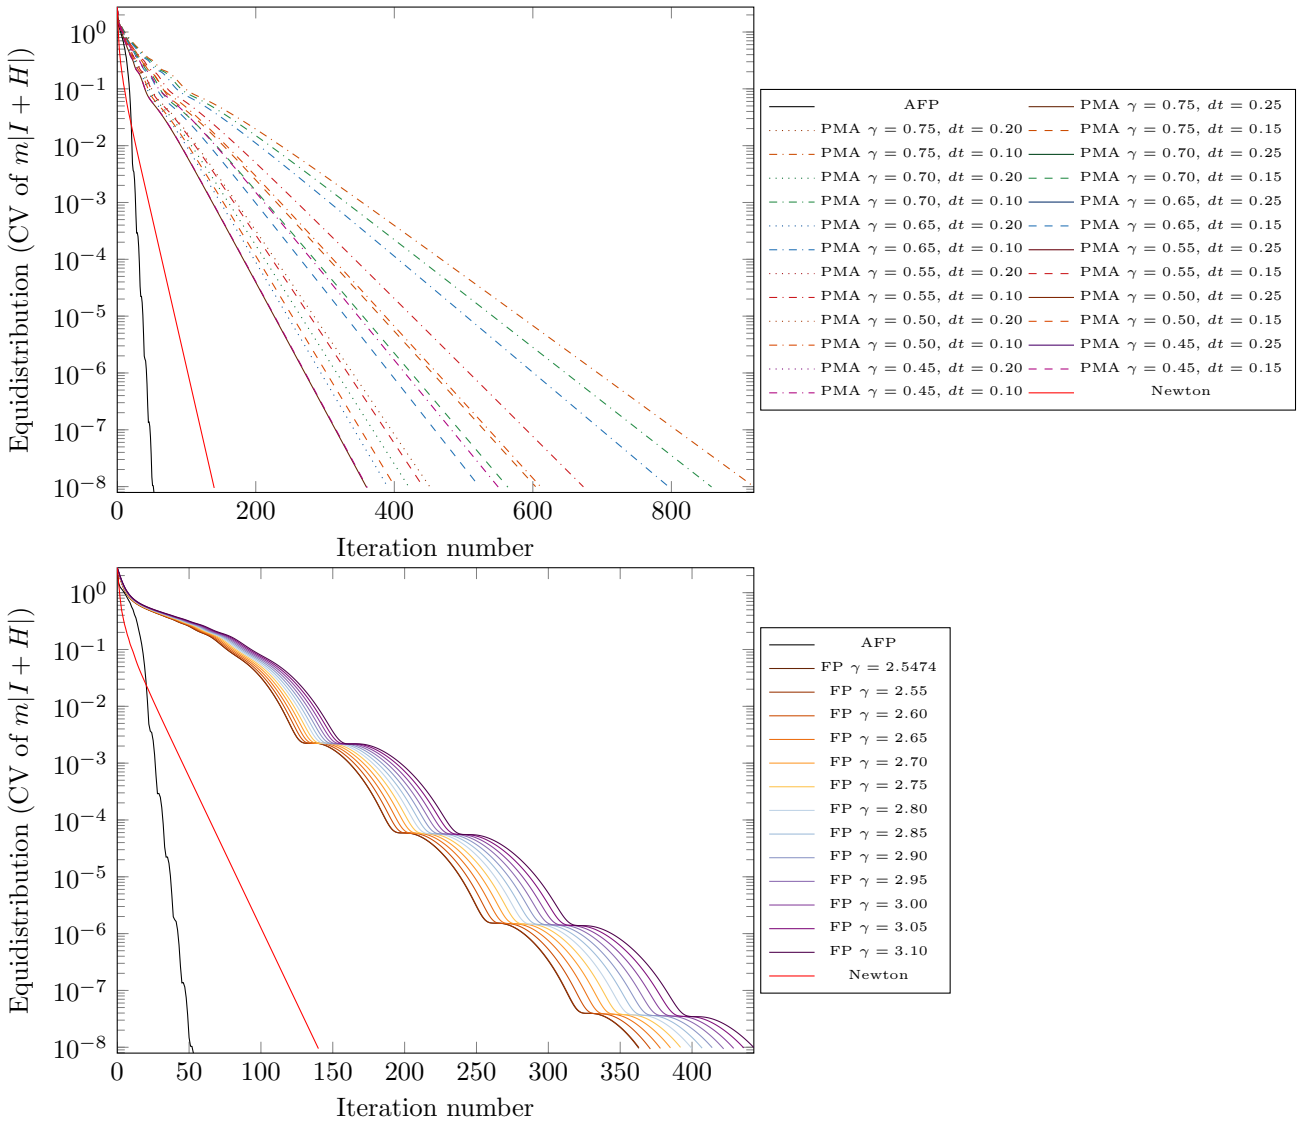

Figure 5: Plots of equidistribution against iteration number for the different methods when applied to the bell test case with an initial 60×60 mesh

4.5 Scaling of the methods with mesh size

(a) Plot of number of iterations taken against mesh size for the ring test case ↓

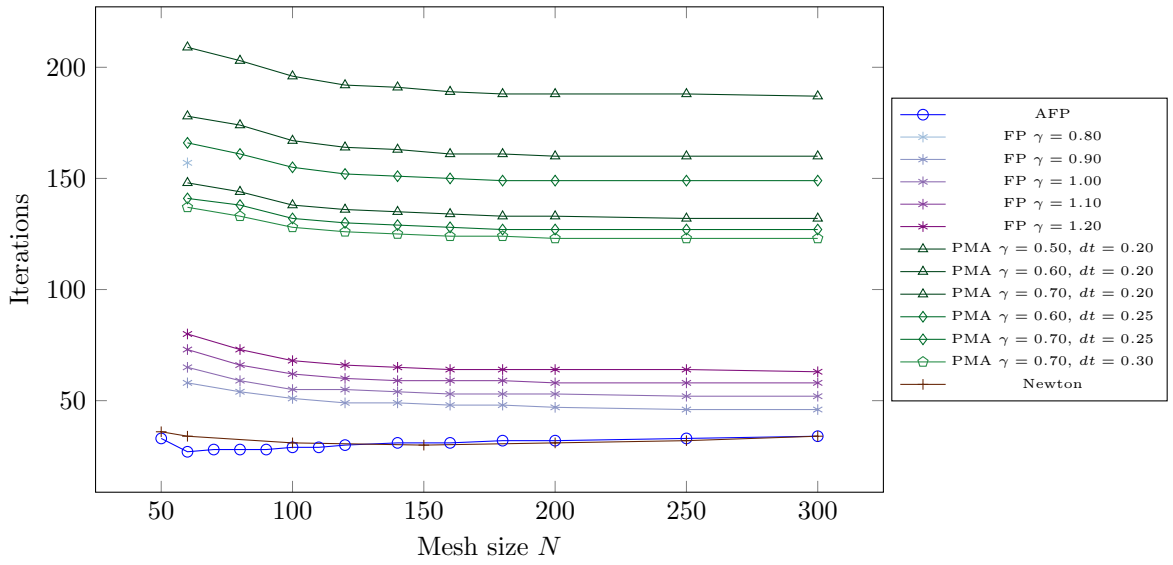

(b) Plot of number of iterations taken against mesh size for the bell test case ↓

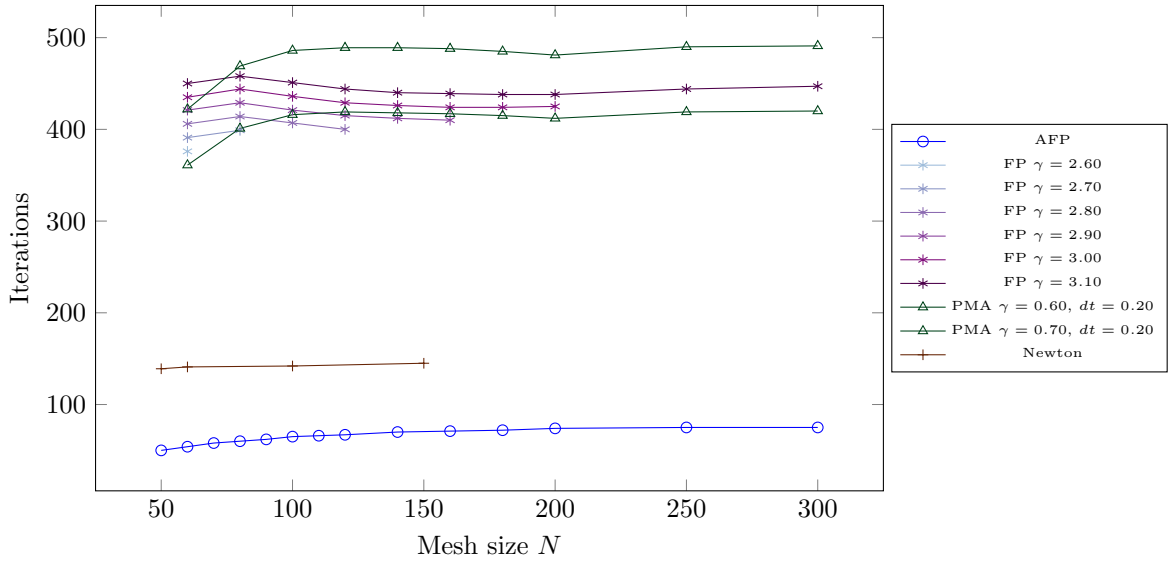

Figure 6: Scaling of the different algorithms as the resolution of the computational mesh increases.

Figure 6 shows how the number of iterations for each method varies as a function of the (square root of the) number of cells in the computational mesh varies. There are multiple lines for both the PMA method and the FP method due to the different free parameters in each technique. In contrast, there is only one line for the AFP method as it has no free parameters. We only show Newton's method for a

single value of the constant δ. Considering first the PMA method, we see that the number of iterations is (for large N) independent of the mesh size. This is consistent with previous studies [Browne et al., 2014]. However, note that for the ring test case the PMA method failed in 3 out of the 9 parameter sets: the lines ($\gamma = 0.5, dt = 0.25$), ($\gamma = 0.5, dt = 0.3$), and ($\gamma = 0.6, dt = 0.3$) all do not appear in Figure 6a. Further, only 2 of the parameter sets converged for the bell test case. This highlights the importance of these free parameters for the tuning of the PMA method.

For the fixed point method, Figure 6 shows that the optimal regularisation parameter γ is strongly problem dependent. For the ring test case, a γ in the interval $[0.8, 1.2]$ was reasonable, whereas in the bell test case $\gamma \in [2.6, 3.1]$ was more appropriate. For the ring test case, Figure 6a shows that the number of iterations is independent of N as in the PMA method. The general pattern that a smaller γ leads to fewer iterations with the FP method can be seen clearly in Figure 6a. There is, however, a limit to how small γ can be before the algorithm fails. This is seen with $\gamma = 0.8$, that only converges for the smallest $N = 60$ case. For the bell test case (Figure 6b) it is clear that the optimal γ is dependent on the mesh size N. The higher the resolution, the larger γ is needed for the method to converge, at the expense of taking more iterations.

For the AFP method, Figure 6 shows that there may be a slight increase in the number of iterations taken as the resolution increases. However, the order of magnitude remains the same and the method converges for every mesh size without having to choose any parameters.

4.6 Problems with the Newton method

Consider now the behaviour of the Newton solver shown in Figure 6. Notice immediately that the Newton method failed for the bell test case when the problem size is greater than 150×150. The method fails due to catastrophic mesh tangling and we have investigated whether this could be caused by numerical errors in the calculation of $\nabla_x \left(\frac{c}{m}\right)$ on the physical mesh.

For the form of the monitor function that we consider in this paper (see (45)) it is a simple exercise to use symbolic algebra to derive an analytic expression for $\nabla_x \left(\frac{c}{m}\right)$. Numerical tests were performed using the analytic rather than the numerical gradient of $\nabla_x \left(\frac{c}{m}\right)$. We found that using the smooth, analytic gradient neither increased the robustness of the Newton method (i.e. achieved convergence when the numerical calculation of the gradient failed) nor increased the efficiency of the method (i.e. the results remained within a single iteration of the numerical gradient calculations). Therefore the slow convergence or divergence of the Newton method appears not to be due to numerical errors in calculating $\nabla_x \left(\frac{c}{m}\right)$

One can also add under-relaxation to the Newton method by increasing the value of δ in (46) – we found that if δ were increased to a level that make the method convergent then the number of iterations needed were prohibitively large. It is therefore not clear why the Newton method fails in some cases or why it is

less convergent than the fixed point method when both are near convergence.

4.7 Use of the adaptive regularisation term

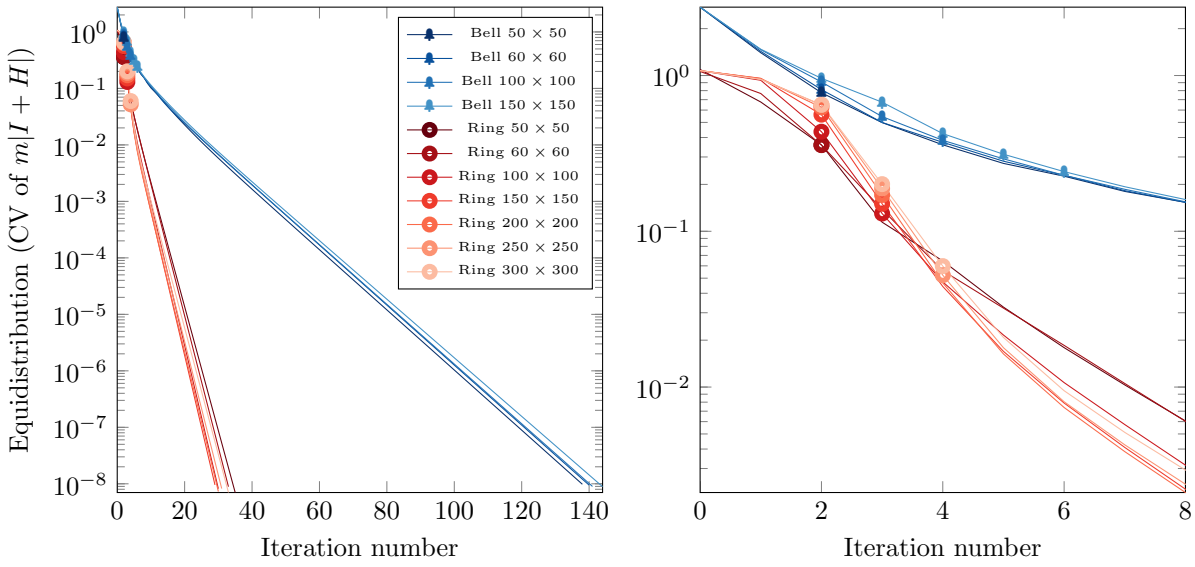

Figure 7: Plots of equidistribution against iteration number for the Newton method with the marks showing the iterations where regularisation occured (i.e. $\gamma \neq 0$ in (22)). On the left is show the convergence for every iteration whilst on the right we zoom in to the initial iterations where the regularisation was needed.

Figure 7 shows the iterations at which the regularisation term in equation (22) was used. In the numerical experiments we have conducted this was only ever needed in the Newton method and not the AFP method. However, when fewer correcting iterations are used to solve the numerical linear algebra problem as discussed in Section 3.6, this regularisation is also needed in the AFP method. Note from Figure 7 that $\gamma \neq 0$ only in the first few iterations, when the solution is far from converged to the unique convex solution.

5 Conclusions

We have introduced a new iterative method for solving the Monge-Ampère equation for mesh redistribution based on linearising the determinant of the Jacobian of the optimal transport map about a previous solution. Using this linearisation we have added a regularisation based on the spectrum of the discrete linearised operator to ensure ellipticity at each step of the iterative procedure. This has resulted in a

method with no free parameters which has been robust to all the test cases we have considered.

We have shown that the method is more efficient than a parabolic relaxation of the Monge-Ampère equation and is significantly more efficient than a fixed point method based on a linearisation about 0 when such a method requires a large amount of under-relaxation to converge.

We have considered a complete linearisation of the Monge-Ampère equation that leads to Newton's method to solve the Monge-Ampère equation. We have shown that this method does not always lead to a reduction in the number of iterations needed to solve the Monge-Ampère equation when compared to the AFP method. Newton's method was found to exhibit only linear convergence, and not quadratic or super-linear convergence that would make it more attractive as a solution algorithm compared with the simpler algorithms that can provide linear convergence with similar or better rates.

This article has considered applying different algorithms directly to given monitor functions without applying any *ad hoc* smoothing (also referred to as filtering) of the monitor function. Such smoothing will necessarily lead to smoother meshes (as the regularity of the mesh is equivalent to the regularity of the monitor function), however it obscures some of the convergence properties of the solution process. The results given in this article may therefore be different when smoothing is present.

There is still an open question of what convergence criteria to use when stopping an algorithm for mesh redistribution. In this paper, we have used a global statistic based on the equidistribution measure defined at the centre of each finite volume cell. If a redistributed mesh is going to be used for the numerical solution of a set of PDEs, equidistribution may not need to be satisfied to a tight tolerance, and the condition that it holds at a specific point in each cell is likely not necessary. Future investigations into optimally transported mesh redistribution techniques may wish to consider stopping criteria that hold only somewhere in a given cell.

Finally, note that the results stated here are valid only in Euclidean geometries (specifically on the plane). When solving mesh redistribution on the sphere or other manifolds, an analogous linearisation of the determinant of the Jacobian of the optimal transport map must take into account of the curvature of the manifold. The linearisation technique given in this article may provide a framework to form such a linearisation.

6 Code availability

The codes used to implement the test problems in this paper, using OpenFOAM, are available from http://researchdata.reading.ac.uk/SOME_URL

Link available following article acceptance

The codes are available directly from github: https://github.com/AtmosFOAM/AMMM/releases/tag/

v0.2

7 Acknowledgements

This work was supported by NERC grant NE/M013693/1 (PAB & HW), EPSRC grant EP/L016613/1 (JP) and EPSRC grant EP/P000835/1 (TP).

References

Anderson, E., Bai, Z., Bischof, C., Blackford, S., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., and Sorensen, D. (1999). *LAPACK Users' guide*, volume 9. Siam.

Barles, G. and Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283.

Benamou, J.-D., Froese, B. D., and Oberman, A. M. (2010). Two numerical methods for the elliptic Monge-Ampère equation. *ESAIM: Mathematical Modelling and Numerical Analysis*, 44(4):737–758.

Boyd, J. (2013). *Chebyshev and Fourier Spectral Methods: Second Revised Edition*. Dover Books on Mathematics. Dover Publications.

Brenier, Y. (1991). Polar Factorization and Monotone Rearrangement of Vector-Valued Functions. *Communications on Pure and Applied Mathematics*, XLIV:375–417.

Browne, P., Budd, C., Piccolo, C., and Cullen, M. (2014). Fast three dimensional r-adaptive mesh redistribution. *Journal of Computational Physics*, 275:174–196.

Budd, C., Cullen, M., and Walsh, E. (2013). Monge-Ampére based moving mesh methods for numerical weather prediction, with applications to the Eady problem. *Journal of Computational Physics*, 236:247–270.

Budd, C., Russell, R., and Walsh, E. (2015). The geometry of r-adaptive meshes generated using optimal transport methods. *Journal of Computational Physics*, 282:113–137.

Budd, C. and Williams, J. (2009). Moving mesh generation using the parabolic Monge-Ampère Equation. *SIAM Journal on Scientific Computing*, 31(5):3438–3465.

Budd, C. J., Huang, W., and Russell, R. D. (2009). Adaptivity with moving grids. *Acta Numerica*, 18:1–131.

Cao, W., Huang, W., and Russell, R. D. (2003). Approaches for generating moving adaptive meshes: Location versus velocity. *Applied Numerical Mathematics*, 47(2):121–138.

Chacón, L., Delzanno, G., and Finn, J. (2011). Robust, multidimensional mesh-motion based on Monge–Kantorovich equidistribution. *Journal of Computational Physics*, 230(1):87–103.

Dean, E. and Glowinski, R. (2006a). An augmented Lagrangian approach to the numerical solution of the Dirichlet problem for the elliptic Monge-Ampère equation in two dimensions. *Electron. Trans. Numer. Anal.*, 22:71–96.

Dean, E. and Glowinski, R. (2006b). Numerical methods for fully nonlinear elliptic equations of the Monge–Ampère type. *Computer Methods in Applied Mechanics and Engineering*, 195(13-16):1344–1386.

Delzanno, G., Chacón, L., Finn, J., Chung, Y., and Lapenta, G. (2008). An optimal robust equidistribution method for two-dimensional grid adaptation based on Monge–Kantorovich optimization. *Journal of Computational Physics*, 227(23):9841–9864.

Feng, X. and Jensen, M. (2016). Convergent semi-Lagrangian methods for the Monge-Ampère equation on unstructured grids. *http://arxiv.org/abs/1602.04758*, pages 1–20.

Feng, X. and Neilan, M. (2009). Vanishing moment method and moment solutions for fully nonlinear second order partial differential equations. *Journal of Scientific Computing*, 38(1):74–98.

Froese, B. D. (2012). *Numerical Methods for the Elliptic Monge-Ampère Equation and Optimal Transport*. Phd, Simon Fraser University.

Froese, B. D. and Oberman, A. M. (2011). Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher. *SIAM Journal on Numerical Analysis*, 49(4):1692–1714.

Guba, O., Taylor, M. A., Ullrich, P. A., Overfelt, J. R., and Levy, M. N. (2014). The spectral element method (SEM) on variable-resolution grids: Evaluating grid sensitivity and resolution-aware numerical viscosity. *Geoscientific Model Development*, 7(6):2803–2816.

Harig, S., Chaeroni, Pranowo, W. S., and Behrens, J. (2008). Tsunami simulations on several scales: Comparison of approaches with unstructured meshes and nested grids. *Ocean Dynamics*, 58(5-6):429–440.

Huang, W. and Russell, R. D. (2011). *Adaptive moving mesh methods*, volume 174. Springer.

Kimura, S., Candy, A. S., Holland, P. R., Piggott, M. D., and Jenkins, A. (2013). Adaptation of an unstructured-mesh, finite-element ocean model to the simulation of ocean circulation beneath ice shelves. *Ocean Modelling*, 67:39–51.

Kühnlein, C., Smolarkiewicz, P. K., and Dörnbrack, A. (2012). Modelling atmospheric flows with adaptive moving meshes. *Journal of Computational Physics*, 231(7):2741–2763.

Lakkis, O. and Pryer, T. (2013). A finite element method for nonlinear elliptic problems. *SIAM Journal on Scientific Computing*, 35(4).

Long, D. and Thuburn, J. (2011). Numerical wave propagation on non-uniform one-dimensional staggered grids. *Journal of Computational Physics*, 230(7):2643–2659.

Neilan, M. (2014). Finite element methods for fully nonlinear second order PDEs based on a discrete Hessian with applications to the Monge-Ampère equation. *Journal of Computational and Applied Mathematics*, 263:351–369.

Oberman, A. M. (2008). Wide stencil finite difference schemes for the elliptic Monge-Ampère equation and functions of the eigenvalues of the Hessian. *Discrete and continuous dynamical systems series B*, 10(1):221–238.

Piani, C., Durran, D., Alexander, M. J., and Holton, J. R. (2000). A Numerical Study of Three-Dimensional Gravity Waves Triggered by Deep Tropical Convection and Their Role in the Dynamics of the QBO. *Journal of the Atmospheric Sciences*, 57(22):3689–3702.

Pryer, T. (2012). Applications of nonvariational finite element methods to Monge–Ampère type equations. *http://arxiv.org/abs/1203.0660*, 0(0):7.

Saumier, L. P., Agueh, M., and Khouider, B. (2015). An efficient numerical algorithm for the L2 optimal transport problem with periodic densities. *IMA Journal of Applied Mathematics (Institute of Mathematics and Its Applications)*, 80(1):135–157.

Sulman, M., Williams, J., and Russell, R. (2011). Optimal mass transport for higher dimensional adaptive grid generation. *Journal of Computational Physics*, 230(9):3302–3330.

Vichnevetsky, R. (1987). Wave propagation and reflection in irregular grids for hyperbolic equations. *Applied Numerical Mathematics*, 3(1-2):133–166.

Wang, Y. (2001). An Explicit Simulation of Tropical Cyclones with a Triply Nested Movable Mesh Primitive Equation Model: TCM3. Part I: Model Description and Control Experiment. *Monthly Weather Review*, 1129(6):1370–1394.

Weller, H. (2014). Non-orthogonal version of the arbitrary polygonal C-grid and a new diamond grid. *Geoscientific Model Development*, 7(3):779–797.

Weller, H., Browne, P., Budd, C., and Cullen, M. (2016). Mesh adaptation on the sphere using optimal transport and the numerical solution of a Monge-Ampère type equation. *Journal of Computational Physics*, 308:102–123.

www.openfoam.org (2015). OpenFOAM.

*Todo list

Link available following article acceptance 24